Using quantum theory to reduce the complexity of input-output processes

Jayne Thompson,[1,*] Andrew J. P. Garner,[1] Vlatko Vedral,[2,1,3] and Mile Gu[4,5,1,†]

[1]*Centre for Quantum Technologies, National University of Singapore, 3 Science Drive 2, 117543, Singapore*
[2]*Atomic and Laser Physics, University of Oxford, Clarendon Laboratory, Parks Road, Oxford, OX1 3PU, United Kingdom*
[3]*Department of Physics, National University of Singapore, 3 Science Drive 2, Singapore 117543*
[4]*School of Physical and Mathematical Sciences, Nanyang Technological University, 639673, Singapore*
[5]*Complexity Institute, Nanyang Technological University, 639673, Singapore*

All natural things process and transform information. They receive environmental information as input, and transform it into appropriate output responses. Much of science is dedicated to building models of such systems – algorithmic abstractions of their input-output behavior that allow us to simulate how such systems can behave in the future, conditioned on what has transpired in the past. Here, we show that classical models cannot avoid inefficiency – storing past information that is unnecessary for correct future simulation. We construct quantum models that mitigate this waste, whenever it is physically possible to do so. This suggests that the complexity of general input-output processes depends fundamentally on what sort of information theory we use to describe them.

arXiv:1601.05420v1 [quant-ph] 20 Jan 2016

Every experiment involves applying actions to some system, and recording its corresponding output response. Both inputs and outputs are recorded as classical bits of information, and the system's operational behaviour can always be regarded as an input-output process that transforms inputs to outputs. Quantitative science aims to capture such behaviour within mathematical models – algorithmic abstractions that can simulate future behaviour based on past observations. Meanwhile realizing models of specifically desired input-output processes allows us to design devices that exhibit identical behaviour. This is essential to the engineering of new technology.

In complexity theory, there is great interest in the simplest models – models that replicate a system's future behaviour while storing the least past information [1, 2]. The motivation is two-fold. Firstly from the rationale of Occam's razor, we should posit no more causes of natural things than are necessary to explain their appearances. Every piece of past information a model requires represents a potential cause of future events, and thus simpler models better isolate the true indicators of future behaviour. The second is practical. As we wish to simulate and engineer systems of increasing complexity, there is a pressing need to find simpler methods that utilize more modest memory requirements.

This motivated systematic methods of constructing such models, first for autonomously evolving systems, and later for systems that can behave differently depending on input. The state of the art are ε-transducers, models of input-output processes that are provably optimal – no other means of modeling a given input-output processes can use less past information [3]. The amount of memory such a transducer requires thus quantifies the minimal causes one needs to postulate to track the process's future behaviour, and presents a natural measure of the process's intrinsic complexity. This heralded new ways to understand diverse systems, ranging from generation of structure in evolutionary dynamics to action-perception cycles underlying artificial agents [4–7]. Yet ε-transducers are classical, their optimality only proven among classical models. Recent research indicates that quantum models can more simply simulate stochastic processes that evolve independently of input [8–10]. Can quantum theory also surpass classical limits in modelling general input-output processes that transform information?

Here, we present systematic methods to construct quantum transducers – quantum models that can be simpler than their classically optimal counterparts. The resulting constructions exhibit significant generality; they can improve upon optimal classical transducers whenever it physically possible to do so. Our work indicates that classical models waste information unavoidably and this waste can be mitigated via quantum processing. As a consequence, existing notions about which input-output processes are complex may fundamentally change in the quantum regime.

Notation and Framework. Our research adopts the tools of computational mechanics [1–3]. We seek to model general systems that at each discrete time-step t can be 'kicked' in a number of different ways denoted by some $x^{(t)}$ selected from a set of possible inputs $\boldsymbol{\mathcal{X}}$. In response, the system emits some $y^{(t)}$ among a set a possible outputs $\boldsymbol{\mathcal{Y}}$.

This framework encompasses autonomously evolving systems as a special case, such that there no choice how the system is kicked at each time-step. The system is then fully characterized by a stochastic process – a sequence of random variables $\overleftrightarrow{Y} = \ldots Y^{(-1)}Y^{(0)}Y^{(1)}\ldots$, where $Y^{(t)}$ governs $y^{(t)}$ for each t. By convention, $t = 0$ represents the present, such that $\overleftarrow{Y} = \ldots Y^{(-2)}Y^{(-1)}$ and $\overrightarrow{Y} = Y^{(0)}Y^{(1)}\ldots$ govern the past and future outputs $\overleftarrow{y} = \ldots y^{(-2)}y^{(-1)}$, and $\overrightarrow{y} = y^{(0)}y^{(1)}\ldots$ respectively.

General input-output processes behave differently depending on input. For each bi-infinite input sequence $\overleftrightarrow{x}$, their output behaves according to a different stochas-

tic process $\overleftrightarrow{Y}|\overleftrightarrow{x}$. The family $\{\overleftrightarrow{Y}|\overleftrightarrow{x}\}_{\overleftrightarrow{x}\in\overleftrightarrow{\boldsymbol{\mathcal{X}}}}$ then completely describes the black-box behaviour of a given input-output process. In particular, if $\overleftrightarrow{x}$ is governed by some stochastic process $\overleftrightarrow{X}$, it will output $\overleftrightarrow{y}$ with probability

$$P[\overleftrightarrow{Y}=\overleftrightarrow{y}]=\sum_{\overleftrightarrow{x}} P[\overleftrightarrow{Y}|\overleftrightarrow{X}=\overleftrightarrow{x}]P[\overleftrightarrow{X}=\overleftrightarrow{x}]. \quad (1)$$

Therefore, an input-output process acts as a transducer on stochastic processes – it transforms one stochastic process (the input) into another (output).

A mathematical model of a given process should be capable of replicating its black-box behaviour. Consider two black boxes, one enclosing the original process and the other enclosing a computer executing a model of the process – such that the only way to interact with the contents of either box is by feeding it inputs and observing the corresponding outputs. The two boxes should be operational indistinguishable; such that they exhibit statistically identical output for any choice of input.

Formally, each instance of an input-output process has a specific past $\overleftarrow{z}=(\overleftarrow{x},\overleftarrow{y})$, an ordered pair of past inputs and outputs (this implicitly assumes causality, such that future inputs cannot retroactively change past outputs). On future input $\overrightarrow{x}$, the process will exhibit a corresponding conditional future behaviour $P[\overrightarrow{Y}|\overrightarrow{X}=\overrightarrow{x},\overleftarrow{Z}=\overleftarrow{z}]$. A model for this process must exhibit identical behaviour. That is, each model records $s(\overleftarrow{z})$ in some physical memory Ξ in place of $\overleftarrow{z}$, such that the model can generate a random variable $\overrightarrow{Y}$ according to $P[\overrightarrow{Y}|\overrightarrow{x},\overleftarrow{z}]$ when given only s (see Fig 1). In complexity theory, there is a great interest in finding the simplest models – models in which the amount of memory required to store s (i.e., the entropy of Ξ) is minimized.

Simplest Classical Models. The *ε-transducer* is the provably simplest classical model [3]. The construction reasons that to simulate future behaviour, we need not differentiate pasts that have statistically identical future input-output relations. Formally, we define the equivalence relation $\overleftarrow{z}\sim_\varepsilon\overleftarrow{z}'$ whenever two pasts, $\overleftarrow{z}$ and $\overleftarrow{z}'$, exhibit statistically coinciding future input-output behaviour, i.e., whenever $P[\overrightarrow{Y}|\overrightarrow{X},\overleftarrow{Z}=\overleftarrow{z}]=P[\overrightarrow{Y}|\overrightarrow{X},\overleftarrow{Z}=\overleftarrow{z}']$. This partitions the space of all pasts into a collection of equivalence classes $\boldsymbol{\mathcal{S}}=\{s_i\}$. In literature, each $s_i\in\boldsymbol{\mathcal{S}}$ is known as a causal state, and ε as the encoding function that maps each past to its corresponding causal state.

When supplied with causal state s_i, the ε-transducer operates according to the transition elements

$$T_{ij}^{y|x}=P[S^{(t)}=s_j,Y^{(t)}=y|S^{(t-1)}=s_i,X^{(t)}=x], \quad (2)$$

the probability a transducer in causal state $s_i\in\boldsymbol{\mathcal{S}}$ will transition to causal state $s_j\in\boldsymbol{\mathcal{S}}$ while emitting $y\in\boldsymbol{\mathcal{Y}}$, conditioned on receiving input $x\in\boldsymbol{\mathcal{X}}$. Iterating through

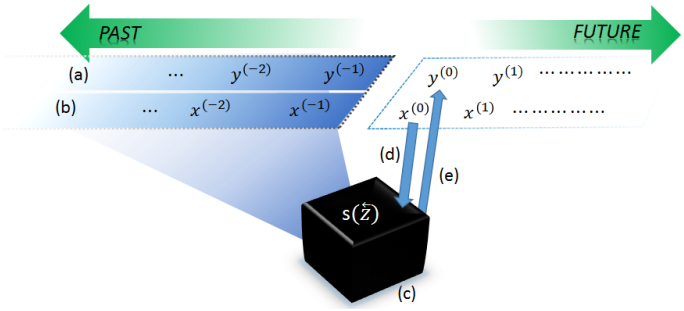

FIG. 1: **Modelling a general input-output process.** Each instance of an input-output process feature some specific sequence of past inputs (a) and past outputs (b). A model of such a process describe a systematic way of storing relevant information within a physical system (c), such that for any future input (d), it can replicate the correct statistical output (e).

this process will generate output behaviour statistically identical to that of the original system we seek to model.

For each stationary input process $\overleftrightarrow{X}$, causal state s_i occurs with probability $p_X(i)=P[\varepsilon(\overleftarrow{z})=s_i]$. The ε-transducer will thus exhibit internal entropy

$$C_X=-\sum_i p_X(i)\log p_X(i). \quad (3)$$

ε-transducers are the provably simplest classical models; any other encoding function $s(\overleftarrow{z})$ that generates correct future statistics exhibits greater entropy. Thus complexity theorists proposed C_X as a quantifier of complexity [3], the rationale being that it characterizes the minimum memory any model must store when simulating $\{\overrightarrow{Y}|\overleftrightarrow{x}\}$ on input $\overleftrightarrow{X}$.

In the special case of modelling autonomous input-output process (i.e., a stochastic process), causal states are reduced to equivalences classes on the set of past outputs $\overleftarrow{\boldsymbol{\mathcal{Y}}}$. $T_{ij}^{y|x}$ become independent of x and are reduced to transition elements T_{ij}^{y}. Here, the ε-transducers are known as ε-machines and C_X as the *statistical complexity*. This measure has been used extensively to quantify the structure of various stochastic processes [11–14].

For general input-output processes, C_X is $\overleftrightarrow{X}$-dependent and is known as the *input-dependent statistical complexity*. In certain pathological cases (e.g. always inputting the same input at every time-step), the transducer may have zero probability of being in particular causal states; potentially leading to drastic reduction on C_X. In general, we consider *non-pathological* inputs, such that the transducer has non-zero probability to be in each causal state, i.e., $p_X(i)>0$ for all i.

It is also often useful to quantify the intrinsic structure of input-output processes without referring to a specific input process [3]. One proposal is $\overline{C}=\sup_X C_X$, which measures how much memory is required in the worst case scenario. $\overline{C}$ is considered an intrinsic measure of

complexity of the system, and is known as the *structural complexity*.

Classical Inefficiency. Even optimal classical models may still store unnecessary information. Consider an ε-transducer equipped with causal states $\boldsymbol{\mathcal{S}}$ and transition elements $T_{ij}^{y|x}$. Suppose there exists $s_i, s_j \in \boldsymbol{\mathcal{S}}$ such that irrespective of input x, both s_i and s_j have non-zero probability of transitioning to some coinciding $s_k \in \boldsymbol{\mathcal{S}}$ while emitting a coinciding output $y \in \boldsymbol{\mathcal{Y}}$ – i.e., both $T_{ik}^{y|x}$ and $T_{jk}^{y|x}$ are non-zero. This implies that at the subsequent time-step, it is impossible to infer which of the two causal states the transducer was previously in with certainty. Thus some of information being stored during the previous time-step must inevitably be lost. We refer to any ε-transducer that exhibits this condition as being *step-wise inefficient*.

Quantum Transducers. Quantum processing can mitigate this inefficiency. Whenever the ε-transducer of a given input-output process is step-wise inefficient, we can construct a quantum transducer that is provably simpler. Our construction assigns each causal state s_i a corresponding quantum causal state

$$|s_i\rangle = \bigotimes_x |s_i^x\rangle, \text{ with } |s_i^x\rangle = \sum_k \sum_y \sqrt{T_{ik}^{y|x}}|y\rangle|k\rangle, \quad (4)$$

where $\bigotimes_x$ represents the direct product over all possible inputs $x \in \boldsymbol{\mathcal{X}}$, while $|y\rangle$ and $|k\rangle$ denote some orthonormal basis for spaces of dimension $|\boldsymbol{\mathcal{Y}}|$ and $|\boldsymbol{\mathcal{S}}|$ respectively. The set of quantum causal states $\{|s_i\rangle\}_i$ then form the state space of the quantum transducer. Note that although each state is represented as a vector of dimension $(|\boldsymbol{\mathcal{Y}}||\boldsymbol{\mathcal{S}}|)^{|\boldsymbol{\mathcal{X}}|}$, in practice they span a Hilbert space of dimension at most $n = |\boldsymbol{\mathcal{S}}|$, and can thus always be stored losslessly within a n-dimensional quantum system (see subsequent example and methods).

For an input process $\overleftrightarrow{X}$, the quantum transducer thus has input-dependent complexity of

$$Q_X = -\mathrm{Tr}\left[\rho_X \log \rho_X\right], \quad (5)$$

where $\rho_X = \sum_i p_X(i)\,|s_i\rangle\langle s_i|$. Note that when there is only one possible input, $|s_i\rangle = |s_i^{(0)}\rangle$. This recovers existing quantum ε-machines that can model autonomously evolving stochastic processes better than their simplest classical counterparts [8]. Our transducers generalize this result to input-output processes. We show that they field the following properties:

1. *Correctness:* For any past $\overleftarrow{z}$ in causal state s_i, a quantum transducer initialized in state $|s_i\rangle$ can exhibit correct future statistical behaviour. i.e., there exists a systematic protocol that, when given $|s_i\rangle$ generates $\overrightarrow{y}$ according to $P[\overrightarrow{Y}|\overrightarrow{X} = \overrightarrow{x}, \overleftarrow{Z} = \overleftarrow{z}]$, for each possible future input $\overrightarrow{x}$.

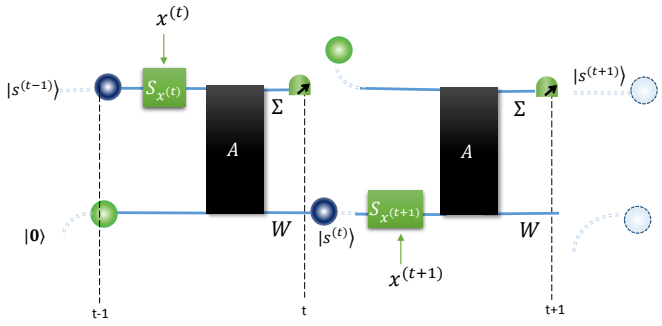

FIG. 2: The quantum circuit which illustrates how a quantum transducer initialized in $|s^{(t-1)}\rangle$, simulates the future behaviour, when supplied with an input sequence $x^{(t)}x^{(t+1)}\ldots$. Upon receiving x at time-step t, it applies a selection operator $S_x : |s_i\rangle \to |s_i^x\rangle$, followed by the quantum operation A that takes each $|s_i^x\rangle$ to some bipartite state $\sum_{y,j}\sqrt{T_{ik}^{y|x}}|y\rangle|s_k\rangle$ with bipartition Σ, spanned by $|y\rangle$, and $\mathcal{W}$, spanned by $|s_k\rangle$ (This is always possible with suitable ancilla, see methods for details.) Σ is emitted as output, while $\mathcal{W}$ is retained as the quantum causal state at the next time-step. Measurement of Σ in the $|y\rangle$ basis by any outside observer yields outcome $y^{(t)}$. The process can then iterated to generate correct outputs at each time-step.

2. *Reduced Complexity:* $Q_X < C_X$ for all non-pathologic input processes $\overleftrightarrow{X}$, whenever the process has a step-wise inefficient ε-transducer.

3. *Generality:* Quantum transducers store less memory whenever it is physically possible to do so. Given an input-output process, either $Q_X < C_X$ for all non-pathological $\overleftrightarrow{X}$, or there exists no physically realizable model that does this.

Correctness guarantees quantum transducers behave statistically identically to their classical counterparts – and are thus operationally indistinguishable from the input-output processes they simulate. The proof is done by explicit construction (see Fig. 2 with details given in methods).

Reduced complexity show step-wise inefficiency is sufficient for quantum transducers to be simpler than their classical counterparts. The proof involves showing that if s_i and s_j have non-zero probability of transitioning to a coinciding causal state s_k while emitting identical y, irrespective of input x, then $|s_i\rangle$ and $|s_j\rangle$ are non-orthogonal (see methods for details). Thus provided two such causal states exist (guaranteed by step-wise inefficiency of the transducer), and each occur with some non-zero probability (guaranteed by non-pathology of the input), $Q_X < C_X$.

Generality establishes that 'step-wise inefficiency' is a necessary condition for any physically realizable quantum model to outperform its classical counterpart. Combined with 'reduced complexity', they imply that 'step-

wise inefficiency' in the sole source of avoidable classical inefficiency and that our particular quantum transducers are general in mitigating this inefficiency. The proof is detailed in the methods, and involves showing that any model which improves upon an ε-transducer that is step-wise efficient allows perfect discrimination of non-orthogonal quantum states.

Together, these three result isolate stepwise inefficiency as the necessary and sufficient condition for quantum models to be simpler than their classical counterparts, and furthermore, establish an explicit construction of such a model. It follows that whenever the ε-transducer is step-wise inefficient, the upper-bound

$$\overline{Q} = \sup_X Q_X, \tag{6}$$

will be strictly less than $\overline{C} = \sup_X C_X$, provided $\sup_X C_X$ is attained for a non-pathological $\overleftrightarrow{X}$. Intuitively, this clause appears natural. If an agent wished to drive a transducer to exhibit the greatest entropy, it is clear that it would be generally advantageous to ensure that the transducer had finite probability of being in each causal state. Nevertheless, as the maximization is highly non-trivial to evaluate, this remains an open conjecture.

Example. We illustrate these results by constructing the quantum transducer of a simple input-output process. Consider a box containing a fair coin. At each time-step an external agent inputs a single bit $x = \{0, 1\}$ by shaking the box with one of two different strengths; causing the coin within to flip with some x-dependent probability. $x = 1$ flips the coin with probability p, while $x = 0$ flips it with probability q, where $0 < p, q < 1$. The new state of the coin, y is then sent out of the box as output. We refer to this process as an *actively perturbed coin.*

First note as $p, q \to 0.5$, the output of the the device becomes completely random, all pasts collapse to a single causal state, and both the quantum and classical structural complexity are trivially zero. In all other cases, the process's future input-output behaviour depends exclusively on the outcome of the previous coin toss, $y^{(-1)}$. Thus the past partitions into two causal states, $s_k = \{\overleftarrow{z} : y^{(-1)} = k\}$, $k = 0, 1$. Future statistics can be generated by transition elements $T_{01}^{1|0} = T_{10}^{0|0} = p$, $T_{00}^{0|0} = T_{11}^{1|0} = 1-p$, $T_{10}^{0|1} = T_{01}^{1|1} = q$, $T_{00}^{0|1} = T_{11}^{1|1} = 1-q$. Consider a simple input process $\overleftrightarrow{X}(u)$ where $x = 1$ is chosen with some fixed probability u at each time-step, the symmetry of the transition elements implies s_0 and s_1 occur with equiprobability. Thus $C_{X(u)} = 1$. Furthermore, as this is the maximum entropy a two-state machine can have, the classical structural complexity of the actively perturbed coin, $\bar{C}$, is also 1.

Such an ε-transducer exhibits step-wise inefficiency. s_0 and s_1 have a finite probability of transitioning to the same causal state while emitting the same output, whether $x = 0$ or $x = 1$. Thus, there exists no systematic method to perfectly retrodict whether an ε-transducer is initialized in s_0 or s_1, regardless of how we choose future inputs. Some information used to perfectly distinguish these two states does not appear to correlate with future black-box behaviour. In contrast the quantum transducer makes use of non-orthogonal causal states,

$$|s_0\rangle = \sqrt{r}|0\rangle + \sqrt{1-r}|1\rangle, \qquad |s_1\rangle = |0\rangle, \tag{7}$$

where $r = 16pq(1-p)(1-q)$ (see methods for details on how future statistics can be generated). Clearly, $|\langle s_0|s_1\rangle| = \sqrt{r} > 0$, provided $0 < p, q < 1$. Thus $Q_X < C_X$ for all input processes $\overleftrightarrow{X}$. Furthermore, the quantum structural complexity $\bar{Q} = \sup_X Q_X$ is attained for any input process where $|s_0\rangle$ and $|s_1\rangle$ occur with equiprobability, and is thus attained for any $\overleftrightarrow{X} = \overleftrightarrow{X}(u)$. The improvement is significant, and can be explicitly evaluated (see Fig. 3).

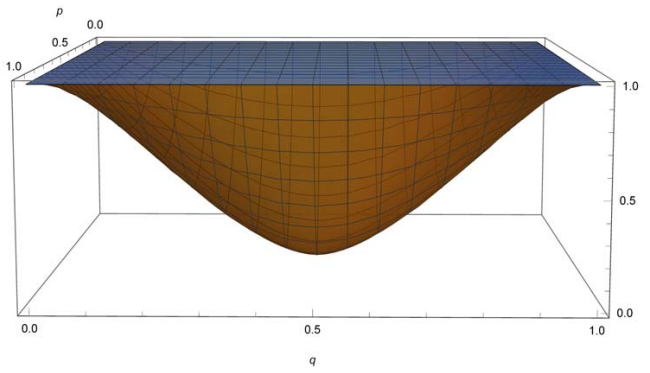

FIG. 3: Plot of the maximum memory required by a quantum transducer $\overline{Q}$ (orange surface) to simulated the actively perturbed coin versus its classical counterpart, the structural complexity $\overline{C}$ (blue surface) for various p and q. While the ε-transducer generally requires 1 bit of memory, the quantum transducer generally requires less, and becomes increasingly more efficient as $p, q \to 0.5$.

Specifically as $p, q \to 0.5$, $Q_X/C_X \to 0$, and the quantum transducer uses a negligible amount of memory compared to its classical counterpart. The intuition is that as we add noise, the future black-box behaviour of a coin initialized in heads becomes increasingly similar to one in tails. A classical transducer nevertheless must distinguish the two states of the coin, and requires a single bit a memory. A quantum transducer, however, has the freedom to only *partially* distinguish the two coin states, and thereby simulate the same behaviour with reduced internal entropy. This mirrors similar behaviour for modelling certain stochastic processes, where the quantum models progressively outperformed their classical counterparts in modelling systems that becomes progressively more random [15]. Our example indicates this phenomena remains when modelling much more general processes that depend on user input.

Discussion. In computational mechanics, ε-transducers are the provably simplest models of input-output processes. Their internal entropy is a quantifier of structure – the rationale being any device capable of replicating the process's behaviour must track at least this much information. Here, we generalize this formalism to the quantum regime. We propose a systematic method to construct quantum transducers that are in general simpler than their simplest classical counterparts; in the sense that quantum transducers store less information whenever it is physically possible to do so. This methodology is applied to explicitly demonstrate quantum models can be far more efficient, even in simulating fairly simple input-output processes. Thus our work indicates the perceived complexity of input-output processes generally depends on what type of information theory we use to describe them.

This connection motivates several avenues of further research, both in exploring the ramifications of measuring structure quantum mechanically, and in motivating new applications for quantum technology. Classical transducers, for example, are used to capture emergence of complexity under evolutionary pressures [16], distil structure within cellular automata [17], and characterize the dynamics of sequential learning [18], and it would be interesting to see how these ideas change in the quantum regime. Meanwhile, recent results suggests the inefficiency of classical models may incur unavoidable thermodynamic costs [19–21]. In reducing this inefficiency, quantum transducers could offer more energetically efficient methods of transforming information. Any such developments could demonstrate the impact of quantum technologies in domains where their use has not been considered before.

Acknowledgements. We thank Suen Whei Yeap, Blake Pollard, Liu Qing, and Yang Chengran for their input and helpful discussions. We are grateful for funding from the John Templeton Foundation Grant 53914 *"Occam's Quantum Mechanical Razor: Can Quantum theory admit the Simplest Understanding of Reality?"*; the Oxford Martin School; the Ministry of Education in Singapore, the Academic Research Fund Tier 3 MOE2012-T3-1-009; the National Research Foundation of Singapore and in particular NRF Award No. NRF–NRFF2016–02.

METHODS.

Definitions. We first introduce extended terminology. Recall $X^{(t)}$,$Y^{(t)}$,$\mathcal{S}^{(t)}$ represent respectively the random variables governing the input, output and causal state at at time t. $X^{(t)}$ (respectively $Y^{(t)}$) takes values $x^{(t)} \in \boldsymbol{\mathcal{X}}$ (respectively $y^{(t)} \in \boldsymbol{\mathcal{Y}}$). $\overleftarrow{X} = X^{(-1)}X^{(-2)}\dots$ and $\overleftarrow{Y} = Y^{(-1)}Y^{(-2)}\dots$ govern past inputs $\overleftarrow{x} \in \overleftarrow{\boldsymbol{\mathcal{X}}}$ and outputs $\overleftarrow{y} \in \overleftarrow{\boldsymbol{\mathcal{Y}}}$ respectively, while $\overrightarrow{X} = X^{(0)}X^{(1)}\dots$ and $\overrightarrow{Y} = Y^{(0)}Y^{(1)}\dots$ govern future inputs $\overrightarrow{x} \in \overrightarrow{\boldsymbol{\mathcal{X}}}$ and outputs $\overrightarrow{y} \in \overrightarrow{\boldsymbol{\mathcal{Y}}}$. Similarly for finite time intervals we define $Y^{(0:t)} = Y^{(0)}\dots Y^{(t)}$ to govern the outputs $y^{(0:t)} = y^{(0)}\dots y^{(t)}$ and $X^{(0:t)} = X^{(0)}\dots X^{(t)}$ to govern the inputs $x^{(0:t)} = x^{(0)}\dots x^{(t)}$. For compactness, we introduce ordered pairs $Z^{(t)} = (X^{(t)}, Y^{(t)})$ which takes values $z^{(t)} = (x^{(t)}, y^{(t)}) \in \boldsymbol{\mathcal{Z}}$, where $\overleftarrow{\boldsymbol{\mathcal{Z}}}$ represents the space of all possible pasts. In analogy, let $\overleftarrow{z} = (\overleftarrow{x}, \overleftarrow{y})$ represents a particular past, $\overleftarrow{\boldsymbol{\mathcal{Z}}}$ be the space of all possible pasts and $z^{(0:t)} = z^{(0)}\dots z^{(t)} \in \boldsymbol{\mathcal{Z}}^{t+1}$ the input-output pairs from time-steps 0 to t. Let $|\boldsymbol{\mathcal{X}}|$ denote the size of the input alphabet, $|\boldsymbol{\mathcal{Y}}|$ the size of output alphabet and $n = |\boldsymbol{\mathcal{S}}|$ the number of causal states. Without loss of generality, we assume the inputs and outputs are labeled numerically, such that $\boldsymbol{\mathcal{X}} = \{x_i\}_0^{|\boldsymbol{\mathcal{X}}|-1}$ and $\boldsymbol{\mathcal{Y}} = \{y_i\}_0^{|\boldsymbol{\mathcal{Y}}|-1}$.

Each instance of such an input-output process $\{\overleftrightarrow{Y}|\overleftrightarrow{x}\}$ exhibits a specific past $\overleftarrow{z} = (\overleftarrow{x}, \overleftarrow{y})$. When supplied $x^{(0)} = x$ as input, it emits $y^{(0)} = 0$ as output with probability $P[Y^{(0)}|X^{(0)} = x, \overleftarrow{Z} = \overleftarrow{z}]$. The system's past then transitions from $\overleftarrow{z}$ to $\overleftarrow{z}' = (\overleftarrow{x}\,x, \overleftarrow{y}\,y) = (\overleftarrow{x}', \overleftarrow{y}')$. This motivates the *propagating functions* $\mu_{x,y}$ on the space of pasts, $\mu_{x,y}(\overleftarrow{z}) = (\overleftarrow{x}\,x, \overleftarrow{y}\,y)$, which characterizes how the past is updated upon observation of (x, y) at each time-step. Iterating through this process for t timesteps gives expected output $y^{(0:t)}$ with probability $P(Y^{(0:t)}|X^{(0:t)} = x^{(0:t)}, \overleftarrow{Z} = \overleftarrow{z})$ upon receipt of input sequence $x^{(0:t)}$. Taking $t \to \infty$ gives the expected future input-output behaviour $P[\overrightarrow{Y}|\overrightarrow{X} = \overrightarrow{x}, \overleftarrow{Z} = \overleftarrow{z}]$ upon future input sequence $\overrightarrow{x} \in \overrightarrow{\boldsymbol{\mathcal{X}}}$.

A general quantum model of the process defines an encoding function $\aleph$ that maps each past $\overleftarrow{z}$ to some state $\rho_{\overleftarrow{z}}$ within some physical system Ξ. The model is correct, if there exists a family of operations $\mathbb{M} = \{\boldsymbol{\mathcal{M}}_x\}_{x\in\boldsymbol{\mathcal{X}}}$ such that application of $\boldsymbol{\mathcal{M}}_x$ on Ξ replicates the behaviour of inputting x. That is, $\boldsymbol{\mathcal{M}}_x$ acting on $\rho_{\overleftarrow{z}}$ should (1) generate output y with probability $P[Y^{(0)} = y|X^{(0)} = x, \overleftarrow{Z} = \overleftarrow{z}]$ and (2) transition Ξ into state $\aleph[\mu_{x,y}(\overleftarrow{z})] = \rho_{\mu_{x,y}(\overleftarrow{z})}$. (1) ensures the model outputs statistically correct $y^{(0)}$, while (2) ensures the model's internal memory is updated to record the event $z^{(0)} = (x^{(0)}, y^{(0)})$, allowing correct prediction of future inputs. Sequential application of $\boldsymbol{\mathcal{M}}_{\overrightarrow{x}} = \boldsymbol{\mathcal{M}}_{x^{(0)}}, \boldsymbol{\mathcal{M}}_{x^{(1)}}, \dots$ then generates output $\overrightarrow{y}$ with probability $P[\overrightarrow{Y}|\overrightarrow{X} = \overrightarrow{x}, \overleftarrow{Z} = \overleftarrow{z}]$. Let $\Omega = \{\rho_{\overleftarrow{z}}\}_{\overleftarrow{z}\in\overleftarrow{\boldsymbol{\mathcal{Z}}}}$ be the image of $\aleph$. We can now define physically realizable models.

Definition 1 (General Quantum Models)**.** *A general quantum model of an input-output process is triple* $\boldsymbol{\mathcal{Q}} = (\aleph, \Omega, \mathbb{M})$.

Each input stochastic process $\overleftrightarrow{X}$ induces a probability distribution $p_{\overleftarrow{z}}$ and the set of pasts. The internal entropy of Ξ,

$$Q_X[\boldsymbol{\mathcal{Q}}] = -\mathrm{Tr}\left(\rho_X \log \rho_X\right), \tag{8}$$

then gives the model's input-dependent complexity. For the quantum transducers described in the main body, $\Omega = \{|s_i\rangle\}_i$ corresponds to the set of quantum causal states, the encoding function $\aleph$ mapped each past $\overleftarrow{z}$ to $|s_i\rangle$, whenever $\overleftarrow{z} \in s_i$ – i.e., whenever the past belongs to causal state $s_i \in \boldsymbol{\mathcal{S}}$.

Another concept we make use of is *input strategies*. Suppose Bob receives a model of some known input-output process, initialized in some (possibly unknown) state $\rho_{\overleftarrow{z}} \in \Omega$. Bob now wants to drive the model to exhibit some particular future behaviour by using an input strategy – a specific algorithm for deciding what input $x^{(t)}$ he will feed it at each specific $t \geq 0$, purely from the model's black-box behaviour.

Definition 2 (Input Strategy)**.** *An input strategy is a family of functions $F = \{f^{(t)} | t \in \mathbb{Z}^+\}$, where $f^{(t)} : \boldsymbol{\mathcal{Z}}^t \to \boldsymbol{\mathcal{X}}$ is a map from the space of pre-existing inputs-outputs onto the input at time t, such that $x^{(t)} = f^{(t)}(z^{(0:t-1)})$. We denote a sequence of future inputs which is determined using $x^{(t)} = f^{(t)}(z^{(0:t-1)})$, as $\overrightarrow{x}_F$.*

In subsequent proofs, we will invoke input strategies on classical ε-transducers. Here, we denote

$$P_{s_i,F}[\overrightarrow{y}] = P[\overrightarrow{Y} = \overrightarrow{y}|, \overrightarrow{X} = \overrightarrow{x}_F, S^{(-1)} = s_i]. \quad (9)$$

as the probability distribution that governs future outputs $\overrightarrow{Y}$, when a input strategy $\overrightarrow{x}_F$ is applied to an ε-transducer initialized in some causal state $s_i \in \boldsymbol{\mathcal{S}}$.

We also make use of the *trace distance*. For any two classical probability distributions $P[\overrightarrow{Y}]$ and $Q[\overrightarrow{Y}]$, their trace distance is given by $D[P,Q] = \frac{1}{2}\sum_{\overrightarrow{y}\in\overrightarrow{\mathcal{Y}}} |P[\overrightarrow{y}] - -Q[\overrightarrow{y}]|$. Similarly for any two quantum states τ and σ the trace distance $D[\tau,\sigma] = \frac{1}{2}\mathrm{Tr}|\tau - \sigma|$.

Proof of Correctness. Here, we prove a quantum transducer can generate correct future statistical behaviour when supplied with a quantum system Ξ initialized in state $\rho_{\overleftarrow{z}} = |s_i\rangle$. That is, there exists a family of quantum process $\mathbb{M} = \{\boldsymbol{\mathcal{M}}_x\}$ whose action on Ξ produces an output y according to probability distribution $P[\overrightarrow{Y} = y|\overrightarrow{X} = x, S^{(-1)} = s_i]$, while transforming Ξ into $\rho_{\mu(\overleftarrow{z},x,y)} = |\varepsilon(\overleftarrow{x}\,x, \overleftarrow{y}\,y)\rangle$.

Proof. Recall that $|s_i\rangle = \bigotimes_x |s_i^x\rangle$ where $|s_i^x\rangle = \sum_{k=0}^{n-1}\sum_{y\in\boldsymbol{\mathcal{Y}}} \sqrt{T_{ik}^{y|x}}|y\rangle|k\rangle$. Let Σ be the $|\boldsymbol{\mathcal{Y}}|$ dimensional Hilbert space spanned by $|y\rangle$ and $\mathcal{K}$ be the n-dimensional Hilbert space spanned by $|k\rangle$. Then each $|s_i^x\rangle$ lies in $\omega = \Sigma \otimes \mathcal{K}$, and each causal state $|s_i\rangle$ lies within $\mathcal{W} = \omega^{\otimes|\boldsymbol{\mathcal{X}}|}$. The set of causal states $\{|s_i\rangle\}_{i=0}^{n-1}$ span some subspace of $\mathcal{W}$ of dimension at most n. This implies that the causal states can be stored within $\mathcal{K}$ without loss, i.e., there exists quantum states $\{|\tau_i\rangle\}_0^{n-1}$ in $\mathcal{K}$ and a unitary process U such that $U|\tau_i\rangle = |s_i\rangle$ for all i (Note, this assumes appending suitable ancillary subspaces to each $|\tau_i\rangle$). An explicit form for $|\tau_i\rangle$ can be systematically constructed through Gram-Schmidt decomposition [22]. We refer to $|\tau_i\rangle$ as *compressed causal states*, and U as the *decompression operator*.

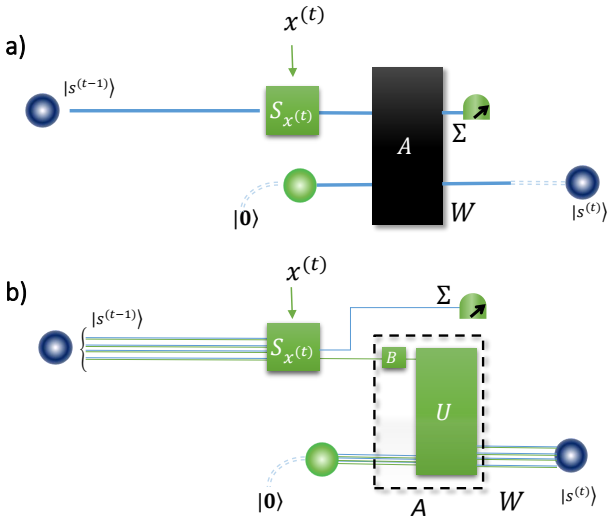

FIG. 4: (a) The process by which the quantum transducer generates future statistics as outlined in Fig. 2. (b) A more detailed breakdown of Fig. (a). Upon input x at time t, the transducer first applies the selection operator S_x. The subsequent operator A can be decomposed two operators, a mapping $B : |s_i^x\rangle \to \sum_{y,k}\sqrt{T_{ik}^{y|x}}|y\rangle|\tau_k\rangle$, and a decompression operator U that rotates each $|\tau_k\rangle$ into $|s_k\rangle$ (always possible when suitable ancillary systems in states $|\mathbf{0}\rangle$ are supplied). Σ is emitted as output, while $\boldsymbol{\mathcal{W}}$ is retained the subsequent causal state $|s^{(t)}\rangle$. This circuit makes it clear that $|\tau_k\rangle$ also make perfectly valid causal states.

We define the *selection operator* $S_x : \mathcal{W} \to \omega$ such that $S_x\left(\bigotimes_{x'}|\phi_{x'}\rangle\right) = |\phi_x\rangle$. Physically, if $\mathcal{W}$ represents a state space of $|\boldsymbol{\mathcal{X}}|$ qudits each with state space ω labelled from 0 to $|\boldsymbol{\mathcal{X}}| - 1$, S_x represents discarding all except the x^{th} qudit. Meanwhile, let B be an quantum operation on ω such that $B|y\rangle|k\rangle = |y\rangle|\tau_k\rangle$. This operation always exists as it can be implemented by Kruass operators $B_{yk} = |y\rangle|k\rangle\langle y|\langle\tau_k|$.

The quantum transducer operates as follows. Upon input x at time t, the transducer applies S_x on Ξ, followed by execution of B. This transforms the state of the system from $|s_i^x\rangle$ to $\sum_{k=0}^{n-1}\sum_{y\in\boldsymbol{\mathcal{Y}}}\sqrt{T_{ik}^{y|x}}|y\rangle|\tau_k\rangle$. Application of the decompression operator then gives $\sum_{k=0}^{n-1}\sum_{y\in\boldsymbol{\mathcal{Y}}}\sqrt{T_{ik}^{y|x}}|y\rangle|s_k\rangle$ on state space $\Sigma\otimes\mathcal{W}$. The machine then emits Σ as the output. Measurement of Σ by an external observer in basis $|y\rangle$ then gives the output y at time t, while inducing $\mathcal{W}$ to transition to the subsequent quantum causal state at time t. $\mathcal{W}$ is then retained in Ξ.

The above procedure establishes a family of quantum operations $\{\mathcal{M}_x\}_{x\in\boldsymbol{\mathcal{X}}}$ that maps each causal state $|s_i\rangle$ to some causal state $|s_j\rangle$ while emitting output y with probability $T_{ij}^{y|x}$. Thus the quantum transducer operates sta-

tistically indistinguishably from its classical counterpart, and must be correct. To simulate the future behaviour of the input-output process subject to a future input sequence $\overrightarrow{x}$, the quantum transducer iterates through the above process by application of $\mathcal{M}_{x^{(0)}}$, $\mathcal{M}_{x^{(1)}}$, etc. in temporal sequence.

Proof of Reduced Complexity and Generality. Here we prove the two remaining results from the main body. Given an arbitrary input-output process, whose ε-transducer and quantum transducer has respect input dependent complexities of C_X and Q_X, we show that

- *Reduced Complexity:* $Q_X \leq C_X$ for all non-pathologic input processes $\overleftrightarrow{X}$, whenever the ε-transducer is step-wise inefficient.
- *Generality:* Given an input-output process, either $Q_X < C_X$ for all non-pathological input processes $\overleftrightarrow{X}$ or there exists no physically realizable model that does this.

Our strategy makes use of the following theorem.

Theorem 1. *For any input-output process the following statements are equivalent:*

(I) $\langle s_i|s_j\rangle = 0$ *for any* $|s_i\rangle, |s_j\rangle \in \{|s_i\rangle\}_i$.
(II) For any pair $(s_i, s_j) \in \boldsymbol{\mathcal{S}} \times \boldsymbol{\mathcal{S}}$, $\exists x \in \boldsymbol{\mathcal{X}}$ *such that* $\forall y \in \boldsymbol{\mathcal{Y}}, s_k \in \boldsymbol{\mathcal{S}}$ *the product* $T^{y|x}_{ik} T^{y|x}_{jk} = 0$.
(III) For any pair $(s_i, s_j) \in \boldsymbol{\mathcal{S}} \times \boldsymbol{\mathcal{S}}$, *there exists an input strategy* F *such that* $D[P_{s_i,F}, P_{s_j,F}] = 1$.
(IV) For any input process $\overleftrightarrow{X}$, *any physically realizable (quantum) model must have complexity at least* C_X.

To prove this theorem, we show **A.** (I) is equivalent to (II). **B.** (II) implies (III), **C.** (III) implies (IV) and **D.** (IV) implies (I).

*Proof of **A.*** We prove this by showing (I) is false iff (II) is false. First, assume (I) is false, such that there exists $|s_i\rangle, |s_j\rangle \in \{|s_i\rangle\}_i$ such that $\langle s_i|s_j\rangle = \Pi_{x\in\boldsymbol{\mathcal{X}}}\left(\sum_{y,k}\sqrt{T^{y|x}_{ik}T^{y|x}_{jk}}\right) > 0$. This implies that for each $x \in \boldsymbol{\mathcal{X}}$ there exists $s_k \in \boldsymbol{\mathcal{S}}, y \in \boldsymbol{\mathcal{Y}}$ such that $T^{y|x}_{ik}T^{y|x}_{jk} \neq 0$. Thus (II) is false. Meanwhile assume (II) is false, i.e., there exists s_i and s_j such that $\forall x$ we can find $y_x \in \boldsymbol{\mathcal{Y}}, s_{k_x} \in \boldsymbol{\mathcal{S}}$ for which $T^{y_x|x}_{ik_x}T^{y_x|x}_{jk_x} > 0$. It follows that $\langle s_i|s_j\rangle > \Pi_x\left(T^{y_x|x}_{ik_x}T^{y_x|x}_{jk_x}\right) > 0$ and (I) is false.

*Proof of **B.*** To prove this, we introduce the the propagator function $g : \boldsymbol{\mathcal{S}} \times \boldsymbol{\mathcal{X}} \times \boldsymbol{\mathcal{Y}} \to \boldsymbol{\mathcal{S}}$, such that $g(s_i, x, y) = s_k$ iff $T^{y|x}_{ik} \neq 0$. Note that g is always a function by joint unifiliarity of the ε-transducer [3]. That is, the triple $(s^{(t-1)}, x^{(t)}, y^{(t)})$ uniquely determines $s^{(t)}$. Furthermore, g is generally many-to-one.

Consider the following game. Alice initializes an ε-transducer in either s_i or s_j and seals it inside a black box. Alice gives this box to Bob and challenges him to infer whether $S^{(-1)} = s_i$ or $S^{(-1)} = s_j$, based purely on the transducers future black-box behaviour. We first prove that if (II) is true, then for each pair (s_i, s_j) there always exists an input strategy F_{ij} that allows Bob can discriminate $S^{(-1)} = s_i$ from $S^{(-1)} = s_j$ to arbitrary accuracy purely from the transducer's output behaviour.

Specifically (II) implies that for all pairs $(s_i, s_j) \in \boldsymbol{\mathcal{S}} \times \boldsymbol{\mathcal{S}}$, there exists some x_{ij} such that $\forall y \in \boldsymbol{\mathcal{Y}}, s_k \in \boldsymbol{\mathcal{S}}$, $T^{y|x_{ij}}_{ik}T^{y|x_{ij}}_{jk} = 0$. At $t = 0$, Bob inputs $X^{(0)} = x_{ij}$. Let $y = Y^{(0)}$ be the corresponding output. There are two possibilities:

(a) One of s_i or s_j never outputs y upon receipt of input x – i.e., either $P[Y^{(0)} = y|X^{(0)} = x_{ij}, S^{-1} = s_i] = 0$ or $P[Y^{(0)} = y|X^{(0)} = x_{ij}, S^{-1} = s_j] = 0$.

(b) $g(s_i, y, x_{ij}) \neq g(s_j, y, x_{ij})$.

If (a) occurs, then Bob can immediately determine whether $S^{(-1)} = s_i$ or $S^{(-1)} = s_j$ and we are done. If (b) occurs, let $s'_i = g(s_i, y, x_{ij})$ be the new causal state if $S^{(-1)} = s_i$. Let $s'_j = g(s_j, y, x_{ij})$ be the new causal state if $S^{(-1)} = s_j$. Given s'_i and s'_j, (II) implies we can find x'_{ij} such that $\forall y \in \boldsymbol{\mathcal{Y}}, s_k \in \boldsymbol{\mathcal{S}}$ the product $T^{y|x'_{ij}}_{i'k}T^{y|x'_{ij}}_{j'k} = 0$. Thus we can repeat the steps above choosing $x^{(1)} = x'_{ij}$. Iterating this procedure defines an input strategy F_{ij}, which determines each input $x^{(t)}$ as a function of observed inputs and outputs. At each point in time t, Bob will be able to identify some $s^{(t)}_i$ which is the current causal state if $S^{(-1)} = s_i$, and some $s^{(t)}_j$ which is the current causal state if $S^{(-1)} = s_j$.

Eventually, either scenario (a) will occur allowing us to perfectly rule out one of $S^{(-1)} = s_i$ or $S^{(-1)} = s_j$. Or alternatively in the limit of an infinite number of time steps, Bob can synchronize the ε-transducer based on the observed inputs and outputs (that is, the causal state at time t is entirely determined by observation of the past in limit of large t [1, 3]). Thus Bob can determine $s^{(t)}$ in the limit as $t \to \infty$, allowing inference of whether $S^{(-1)} = s_i$ or $S^{(-1)} = s_j$. This constitutes an explicit input strategy F_{ij} that allows Bob to discriminate between $S^{(-1)} = s_i$ and $S^{(-1)} = s_j$ based on future input-output behaviour $\overrightarrow{y}$ to any arbitrary accuracy. Accordingly $D[P_{s_iF_{ij}}, P_{s_j,F_{ij}}] = 1$.

*Proof of **C.*** We prove this via its contrapositive. That is, suppose (IV) is false, such there exists quantum model $\boldsymbol{\mathcal{Q}}'$ with identical input-output relations to the process' ε-transducer, and that this model stores $Q_X[\boldsymbol{\mathcal{Q}}'] < C_X$ for some $\overleftrightarrow{X}$. We will show that if (III) is true then we will violate the data processing inequality [23].

We first make use of the following observation: If a model $\boldsymbol{\mathcal{Q}}' = (\aleph', \Omega', \mathbb{M}')$ satisfies $\aleph'(\overleftarrow{z}) \neq \aleph'(\overleftarrow{z}')$ for some $\overleftarrow{z} \sim_\varepsilon \overleftarrow{z}'$, then we can always construct an alternative model $\boldsymbol{\mathcal{Q}} = (\aleph, \Omega, \mathbb{M})$ such that $Q_X[\mathcal{Q}] \leq Q_X[\mathcal{Q}']$ for all input processes $\overleftrightarrow{X}$, and $\aleph(\overleftarrow{z}) = \aleph(\overleftarrow{z}') \Leftrightarrow \varepsilon(\overleftarrow{z}) =$

$\varepsilon(\overleftarrow{z}')$ for all $\overleftarrow{z}, \overleftarrow{z}' \in \overleftarrow{\boldsymbol{\mathcal{Z}}}$ (This is a consequence of the concavity of entropy, see methods in [15]). i.e., for any model $\boldsymbol{\mathcal{Q}}'$ with quantum states Ω' not in 1-1 correspondence with causal states, there always exists a simpler model $\boldsymbol{\mathcal{Q}}$ whose quantum states are in 1-1 correspondence with causal states.

Thus, falsehood of (IV) implies there must exist some quantum model $\boldsymbol{\mathcal{Q}} = (\aleph, \Omega, \mathbb{M})$ such that (i) $\aleph(\overleftarrow{z}) = \aleph(\overleftarrow{z}')$ if and only if $\overleftarrow{z} \sim_\varepsilon \overleftarrow{z}'$ and (ii) $Q_X[\boldsymbol{\mathcal{Q}}] < C_X$ for some $\overleftrightarrow{X}$. Now by virtue (ii), there must exist two states $\rho_i, \rho_j \in \Omega$ such that the trace distance

$$D[\rho_i, \rho_j] < 1. \tag{10}$$

The data processing inequality therefore implies that any quantum operation $\boldsymbol{\mathcal{M}}_{\overrightarrow{x}} : \Xi \rightarrow \overrightarrow{\boldsymbol{\mathcal{Y}}}$ which generates future output statistics must satisfy $D[\boldsymbol{\mathcal{M}}_{\overrightarrow{x}}(\rho_i), \boldsymbol{\mathcal{M}}_{\overrightarrow{x}}(\rho_j)] \leq D[\rho_i, \rho_j] < 1$.

However, all models of the same input-output process have identical black-box behaviour. In particular the ε-transducer of the input-output process that $\boldsymbol{\mathcal{Q}}$ models, must behave identically to $\boldsymbol{\mathcal{Q}}$. As such, there exists two causal states of the classical ε-transducer, $s_i, s_j \in \boldsymbol{\mathcal{S}}$ such that $D[P_{s_i,F}, P_{s_j,F}] = D[\boldsymbol{\mathcal{M}}_{\overrightarrow{x}_F}(\rho_i), \boldsymbol{\mathcal{M}}_{\overrightarrow{x}_F}(\rho_j)] < 1$ for all possible input strategies F. This directly contradicts (III), and thus (III) implies (IV).

Proof of ***D.*** Observe that the quantum transducer is a physically realizable model. Thus (IV) implies that $Q_X \geq C_X$ for all $\overleftrightarrow{X}$. However, we note from our construction $Q_X \leq C_X$ for all $\overleftrightarrow{X}$. Therefore $Q_X = C_X$. Noting that the quantum causal states for the quantum transducers are all pure (i.e., $\Omega = \{|s_i\rangle\}$, this implies all $|s_i\rangle$ are mutually orthogonal.

Proof of Main Result. Reduced complexity and generality can be reduced to consequences of the above theorem. Specifically given a particular input-output process, falsehood of (II) implies that its transducer is stepwise inefficient. Meanwhile, falsehood of (I) implies that $Q_X < C_X$ for all non-pathological $\overleftrightarrow{X}$. Thus reduced complexity is implied by equivalence of (I) and (II). Generality is proven by contradiction. Assume that for some non-pathological $\overleftrightarrow{X}$, quantum transducers yield no improvement (i.e. $Q_X = C_X$) but some other physically realizable model has complexity less than C_X. The former implies (I) is true, the latter implies (IV) is false, which violates the theorem above. Thus, both reduced complexity and generality must hold.

Simulating the actively perturbed coin. Here, we detail how the two quantum causal states in Eq. (7) allows simulation of the actively perturbed coin.

Suppose that at $t = 0$, the quantum transducer is in state $|s_j\rangle$, and receives input x. To generate output y, it first initializes three ancilla qubits, each in state $|0\rangle$, to create the joint four qubit state $|s_j 000\rangle$. We now apply a 4-qubit unitary U such that $U|s_j\rangle|000\rangle = |s'_j\rangle$ for $j = 0, 1$, where

$$|s'_0\rangle = |\phi_p\rangle \otimes |\phi_q\rangle, \qquad |s'_1\rangle = \hat{X}_1\hat{X}_2\hat{X}_3\hat{X}_4|s'_0\rangle.$$

Here, $|\phi_p\rangle = \sqrt{1-p}|00\rangle + \sqrt{p}|11\rangle$ and $|\phi_q\rangle = \sqrt{1-q}|00\rangle + \sqrt{q}|11\rangle$, and $\hat{X}_i$ represents the Pauli X operator on the i^{th} qubit. Such a unitary always exists, as $\langle s'_0|s'_1\rangle = \langle s_0|s_1\rangle$.

The transducer then applies the selection operator S_x, which either discards qubits B_1 and B_2 if $x = 0$, or it discards qubits B_3 and B_4 if $x = 1$. Rename the two remaining qubits C_1 and C_2. C_1 and C_2 is then subjected to a unitary V that maps $|00\rangle$ to $|0s_0\rangle$ and $|11\rangle$ to $|1s_1\rangle$ (this is always possible as $\langle 0s_0|1s_1\rangle = 0$). The transducer emits qubit C_1 as output and retains qubit C_2 as the causal state at the subsequent time-step. Measurement of C_1 in the computational basis by an outside observer then yields y. Iteration of this procedure replicates correct future input-output statistics.

* Electronic address: thompson.jayne2@gmail.com
† Electronic address: cqtmileg@nus.edu.sg

[1] C. R. Shalizi and J. P. Crutchfield, Journal of statistical physics **104**, 817 (2001).
[2] J. P. Crutchfield and K. Young, Physical Review Letters **63**, 105 (1989).
[3] N. Barnett and J. P. Crutchfield, [arXiv:1412.2690] (2014).
[4] R. Meddis, The Journal of the Acoustical Society of America **83**, 1056 (1988).
[5] F. Rieke, *Spikes: exploring the neural code* (MIT press, 1999).
[6] N. Tishby and D. Polani, in *Perception-action cycle* (Springer, 2011), pp. 601–636.
[7] G. Gordon, D. M. Kaplan, B. Lankow, D. Y.-J. Little, J. Sherwin, B. A. Suter, and L. Thaler, Frontiers in systems neuroscience **5** (2011).
[8] M. Gu, K. Wiesner, E. Rieper, and V. Vedral, Nature communications **3**, 762 (2012).
[9] J. R. Mahoney, C. Aghamohammadi, and J. P. Crutchfield, [arXiv:1508.02760] (2015).
[10] R. Tan, D. R. Terno, J. Thompson, V. Vedral, and M. Gu, The European Physical Journal Plus **129**, 1 (2014).
[11] P. Tiňo and M. Köteles, Neural Networks, IEEE Transactions on **10**, 284 (1999).
[12] H. Larrondo, C. González, M. Martin, A. Plastino, and O. Rosso, Physica A: Statistical Mechanics and its Applications **356**, 133 (2005).
[13] W. Gonçalves, R. Pinto, J. Sartorelli, and M. de Oliveira, Physica A: Statistical Mechanics and its Applications **257**, 385 (1998).
[14] J. B. Park, J. W. Lee, J.-S. Yang, H.-H. Jo, and H.-T. Moon, Physica A: Statistical Mechanics and its Applications **379**, 179 (2007).
[15] W. Y. Suen, J. Thompson, A. J. P. Garner, V. Vedral, and M. Gu, [arXiv:1511.05738] (2015).
[16] J. P. Crutchfield and O. Görnerup, Journal of The Royal Society Interface **3**, 345 (2006).

[17] J. E. Hanson and J. P. Crutchfield, Journal of Statistical Physics **66**, 1415 (1992).
[18] J. P. Crutchfield and S. Whalen, PLoS Computational Biology, **8**, e1002510 (2012).
[19] K. Wiesner, M. Gu, E. Rieper, and V. Vedral, in *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences* (The Royal Society, 2012), vol. 468, pp. 4058–4066.
[20] A. J. P. Garner, J. Thompson, V. Vedral, and M. Gu, [arXiv:1510.00010] (2015).
[21] A. Cabello, M. Gu, O. Gühne, J.-Å. Larsson, and K. Wiesner, [arXiv:1509.03641] (2015).
[22] B. Noble and J. W. Daniel, *Applied linear algebra*, vol. 3 (Prentice-Hall New Jersey, 1988).
[23] M. A. Nielsen and I. L. Chuang, *Quantum computation and quantum information* (Cambridge university press, 2010).